# *Introductory Guide to the Aeronautical Systems Center DoD High Performance Computing Major Shared Resource Center*

September 15, 2004

## Release Notes

Document last reviewed on January 14, 2004

Issued April 8, 1997.
Initial release.

Issued August 15, 1997.
Update to directory structure and modified cover.

Issued December 11, 1997.
Format changes made to standardize all user guides.

Issued February 10, 1999.
Kerberos information added. Updated to reflect hardware upgrades.

Issued February 29, 2000.
Updated to reflect hardware upgrades.

Issued October 4, 2000.
Updated to reflect system upgrades.

Issued January 15, 2001.
Updated to reflect current hardware configuration.

Issued October 5, 2001.
Updated to reflect current hardware configuration.

Issued August 23, 2002.
Removed outdated information.

Issued June, 2003.
Removed outdated information.

Issued December 17, 2003
Removed references to the GS-320, added SGI Origin 3900 information and updated policy changes.

Issued January 14, 2004
Updated system references to match HPCMO list of common platform names.

Issued September 15, 2004
Updated to reflect IBM P3 decommissioning.

# *Table of Contents*

# *Table of Contents*

## Purpose

Provides Aeronautical Systems Center (ASC) Major Shared Resource Center (MSRC) background information to users and their Service/Agency Approving Authority (S/AAA). ASC MSRC information is presented in sufficient detail to allow individuals to make informed decisions about the services provided by the ASC MSRC. The information is provided to help users and their S/AAA decide which of the High Performance Computing (HPC) MSRCs is most appropriate for their needs.

## Scope

This document provides technical and policy information about the ASC MSRC. It is not a User's Guide. For detailed execution procedures, please refer to the various ASC MSRC HPC systems.

NOTE: System configuration and policies related to HPC platforms are subject to change. Please refer to the ASC MSRC website for the latest information.

http://www.asc.hpc.mil

## A. Overview and Supported CTAs

This document is intended to provide an overview and introduction to using the ASC MSRC, a part of the Department of Defense (DoD) HPC community. The ASC MSRC is physically located in Area B, Building 676, of Wright-Patterson Air Force Base (WPAFB), near Dayton, Ohio.

Its mission is to provide state-of-the-practice HPC capabilities to all Defense Research, Development, Test and Evaluation user communities with special emphasis in the following Computational Technology Areas (CTAs):

- CSM-Computational Structural Mechanics
- CFD-Computational Fluid Dynamics
- CCM-Computational Chemistry and Materials Science
- CEA-Computational Electromagnetics and Acoustics
- CEN-Computational Electronics and Nano-Electronics

The ASC MSRC is comprised of commercial and "in-house" developed scientific and engineering application software, combined with the hardware needed to support those applications. It encompasses more than merely having a collection of software and computational resources; it also includes methodologies to configure and use these resources in a simple and standardized manner.

The ASC MSRC contains a variety of application and file servers. It supports X-windowing devices to facilitate graphical user interfaces (GUIs) in accessing the available computational resources and to perform Scientific Visualization (SciVis). A vast 500-terabyte tape archival storage/retrieval capability is available to ensure the longevity and safety of every customer's important scientific data.

### Customer Support

The ASC MSRC Customer Service Center is available to respond to questions and comments related to the ASC MSRC, from 0700 to 1700 ET Monday through Friday. The Customer Service Center provides a source of real-time assistance for most problems. After hours support is available Monday through Friday, from 5 pm, to 7 pm Eastern Time, and 24 hours on weekends and holidays (voice mail is used only when personnel are not immediately available).

The support analysts will try to help you with anything related to the ASC MSRC: third-party software, UNIX and its variants, the different application servers, etc. If you have any questions about the ASC MSRC, please call the Customer Assistance and Technology Center first. ***If your problem or question is beyond the scope of their expertise, they will refer you to the appropriate resource to resolve it.***

**CONTACT INFORMATION:**

| | |
|---|---|
| **Toll Free Phone:** | **1 888 MSRC ASC**<br>**1-888-677-2272** |
| **Commercial Phone:** | **937-255-0194** |
| **DSN:** | **785-0194** |
| **FAX:** | **937-656-9538** |
| **Email:** | **msrchelp@asc.hpc.mil** |

**Training**

The ASC MSRC also supports an extensive training schedule. Most training is conducted at the ASC MSRC training facility. Training in your facility and/or specialized training courses are available upon request. Please contact the Service Center (see above) for specific training requests. The training schedule is updated regularly on the ASC MSRC website. (http://www.asc.hpc.mil/education/training)

**Visitor Information**

The ASC MSRC is located at WPAFB in Southwest Ohio, about 50 miles north of Cincinnati. Maps and detailed directions from the Dayton, Ohio airport to the facility are available at http://www.asc.hpc.mil/visitors/maps.php. Visitors to the ASC MSRC facility must check in at the Springfield Street gate Welcome Center (Bldg. 58) in Area B, prior to arriving at the ASC MSRC in Building 676. All foreign national visitors should check with the Foreign Disclosure Office five days prior to coming to ASC. Security clearances for official visitors will be handled in the following manner:

Visit requests should be addressed to: ASC/HP
2435 Fifth Street
WPAFB OH 45433-5706
Voice: (937) 255-1419
FAX: (937) 904-5505

## B. Hardware, Network, and Software

The ASC MSRC is comprised of high performance application servers, high availability data servers, a SciVis facility, and an array of applications supporting five CTAs. (see Figure 1). All servers are linked together using Ethernet and Asynchronous Transfer Mode (ATM), interfaces to maximize network speed, throughput, and access. *All of these networks are interconnected and are connected to the Internet*.

A number of scientific and engineering applications are available at the ASC MSRC. These applications are diverse in nature in order to meet the varying needs of the engineering and scientific community; applications for analysis, mathematics, visualization, CAD/CAE, and program development are available on the various MSRC servers

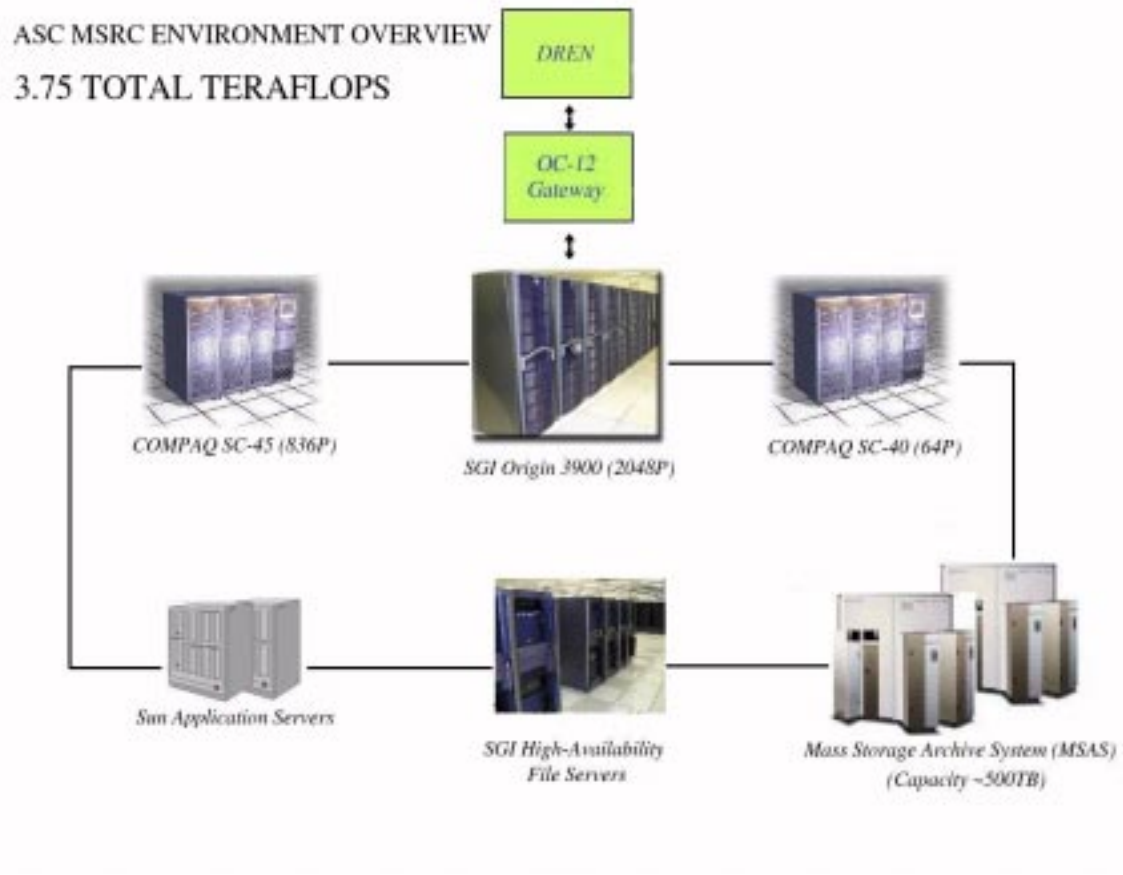


Figure 1: Stylized View of the ASC MSRC

### SGI Origin 3900

The SGI Origin 3900 has a total of 2048 CPUs. Each CPU is a 700 MHz MIPS R14000 processor with a peak speed of 1.2 GFLOPS, providing a total capacity of approximately 2.5 TFLOPs. Each CPU has a primary data cache of 32 KB, a primary instruction cache of 32 KB, and an on-board cache of 8 MB. At this time, the SGI Origin 3900 is divided into 4 separate nodes. The system has a Symmetric MultiProcessing (SMP) architecture with 512 CPUs and 512 GB memory per node.

**COMPAQ SC-40**

The SC-40 has a total of 64 CPUs. Each CPU is a 833 MHz EV6.7 processor with a peak speed of 1 GFLOP, providing a total capacity of approximately 64 GFLOPs. Each CPU has a primary data cache of 64 KB, a primary instruction cache of 64 KB, and an on-board cache of 8 MB. At this time, the SC-40 is divided into 16 separate systems or nodes. The SC-40 has a SMP architecture with 4 CPUs and 4 GB memory per node.

**COMPAQ SC-45**

The SC-45 has a total of 836 CPUs. Each CPU is a 1000 MHz EV6.8 processor with a peak speed of 1 GFLOP, providing a total capacity of approximately 64 GFLOPs. Each CPU has a primary data cache of 64 KB, a primary instruction cache of 64 KB, and an on-board cache of 8 MB. At this time, the SC-45 is divided into 2 systems with 128 nodes with 4 processors per node, totaling 512 processors on one system and 81 nodes with 4 processors per node on the other. The SC-45 has a SMP architecture with 4 CPUs and 4 GB memory per node.

**Archival Storage**

The archival storage system for the ASC MSRC is currently configured with a total capacity of 500 terabytes of near-on-line storage. The system consists of a set of dual Sun E10000 cluster servers with 1.8 terabytes of on-line RAID storage attached to dual-combined StorageTek 9310 tape silos. This facility provides long-term storage for user data files.

**Scientific Visualization**

SciVis resources of the ASC MSRC have been enhanced with the addition of a wide variety of graphic workstations and software. Only unclassified visualization is supported. The equipment is designed to provide high quality graphics in a tightly integrated environment of graphics and video. This equipment includes SGI O2 and Infinite Reality workstations; and various peripherals.

**Network Connectivity**

The ASC MSRC is a critical node of the Defense Research and Engineering Network (DREN), which provides a 622Mbps wide area network service to the MSRC. The internal MSRC network provides 622 Mbps and 155 Mbps ATM, and 100 Mbps fast Ethernet and 10 Mbps Ethernet connectivity between the computational platforms. The following are the Primary Internet Protocol (IP) addresses for the ASC MSRC HPC systems.

**Table 1: MSRC Servers and Addresses**

| System Name | Description | Operating System |
|---|---|---|
| svw10.asc.hpc.mil<br>svw11.asc.hpc.mil | SGI Onyx 3<br>SGI Onyx 2 | IRIX64 6.5 |
| hpc05.asc.hpc.mil<br>hpc09.asc.hpc.mil<br>hpc10.asc.hpc.mil | COMPAQ SC-40<br>COMPAQ SC-45 | TRU64<br>UNIX V5 |
| hpc11.asc.hpc.mil | SGI Origin 3900 | IRIX64 6.5 |

**Software Environment**

All ASC MSRC systems run derivatives of the UNIX System V operating system with vendor-specific enhancements. A large variety of compiler environments, math libraries, programming tools, and third-party analysis applications in support of the CTAs emphasized at ASC are available on the MSRC systems.

**Available Software**

A current list of software available for each system is available at http://www.asc.hpc.mil/software/.

Each third party software package is assigned to an Application Manager, who is responsible for installing, maintaining, and having a general knowledge of the package and its vendor. It is the Application Manager's responsibility to maintain contact with the vendor to ensure that the software package is maintained at the current release level being supported by its vendor, and to work with Customer Assistance and Technology Center in resolving user and application difficulties. The Application Manager is also responsible for tracking application usage and maintaining an adequate number of application licenses.

Supporting the Application Manager is the Systems Administration Staff, which coordinates installation of the system software into the ASC MSRC and ensures the commercial applications software packages are smoothly integrated with the system software.

Currently, the ASC MSRC includes third party applications in the following categories:

- Analysis Software
- Graphics Libraries
- Program Development Environments
- Simulation Software
- Mathematical Software
- Visualization
- Technical Publishing

## C. Permanent/Temporary Storage Environment and Policies

### ASC MSRC Data Concept

The ASC MSRC manages data via a Mass Storage Archival Server (MSAS), a very large temporary shared disk system on each platform called `/workspace` and a personal home directory. Utilities are provided that easily move data between the MSAS, `/workspace` and the personal home directory areas. For complete information regarding use of permanent file storage, refer to the ASC MSRC User's Guide or Archival Storage User's Guide.

### Permanent File Storage

Users are allocated a home directory (referenced locally with the $HOME environment variable) on the SGI High Availability Files Server (HAFS). The directories of the HAFS are automounted on all the HPC computer systems. There is a limit of 450 GB available to the users. Each user is limited to 512 MB with permanent non-migrated storage.

### Mass Storage Archival Server

The archival server uses the Storage Technology Corporation Application Storage Manager to periodically migrate files to tape and to copy files back to disk when a user attempts to access the migrated files. The file status information of the files that are migrated to tape remains on disk which gives the appearance that the files are still stored on disk.

## D. Job Scheduling/Queuing Environment and Policies

### Determining the Correct HPC System

At least four yardsticks exist for measuring job size: CPU processing time, required memory, number of processors, and processor speed. These metrics are not mutually exclusive, e.g., increase the number of processors by parallellizing the code and the processing time will change (hopefully decrease). However, there is not necessarily a direct correlation between the size of the change and the other metric, i.e., doubling the processors will not half the processing time. The ASC MSRC uses memory and number-of-processors standards for estimating application sizes. Using memory size as the yardstick, the COMPAQ SC-40 appears to be a low-end machine. It has relatively small memory, but it has a fast architecture. This general guideline is provided to help in deciding the appropriate platform for your ASC MSRC application. Workstations should be considered for small single processor applications (< 64 megabytes).

The COMPAQ SC-45 is configured and tuned for large to very large parallel applications (1 GB< x < 32 GB) and long-running batch work (8 hrs < x < 24 hrs).

Permissible exceptions are applications developed for a specific architecture. Those applications should be run on the applicable architecture.

**Table 2: System Comparison Matrix**

| System | Architecture | # of Compute Processors | Memory |
|---|---|---|---|
| COMPAQ SC-40 | Parallel | 64 | 64 GB |
| COMPAQ SC-45 | Parallel | 836 | 836 GB |
| SGI Origin 3900 | Parallel | 2048 | 2048 GB |

**Processing Environment Overview and Philosophy**

The ASC MSRC provides both an interactive and batch submission environments. The batch environment is the primary environment for most user work..

**Table 3: Batch Queue Environments**

| System | Batch Queue Environment Server |
|---|---|
| COMPAQ SC-40 | LSF |
| COMPAQ SC-45 | LSF |
| SGI Origin 3900 | LSF |

The BQEs permit users to submit, monitor, and terminate their own unattended batch jobs. This capability is intended for jobs requiring large amounts of memory and/or CPU time that generally run for many hours or days, making them poor candidates for interactive execution. Through the batch queue environments, the user submits a request, or job, which is a standard shell script that the BQE batch monitor executes. This batch shell script may also contain embedded control statements, that are BQE-specific and specify resource requirements, (e.g., CPU time and memory) or runtime parameters (e.g., output file redirection) for the batch job to be executed.

The BQE batch facility presents users with an unattended batch alternative to job execution in a standard UNIX environment via interactive or background execution. BQE is particularly effective in HPC environments where job scheduling becomes a critical factor in maintaining high use of expensive computer systems without permitting them to become grossly oversubscribed and unresponsive.

All ASC MSRC systems are in tremendous demand. They are servicing approximately six hundred users in all three Services and other Defense Agencies. To satisfy our large customer base, striking a balance between the following is necessary:

- The number of concurrent executing batch jobs.
- The desire to keep the costly resources fully used.
- The desire to maintain a reasonable interactive response.
- The ability to schedule batch jobs preferentially that have real “HPC” requirements such as multi-processor jobs, large memory jobs, or jobs with large/special I/O requirements.
- The requirement to launch batch jobs into execution within a reasonable time and get them out of execution within a reasonable time.

Users expect all of the above. Unfortunately, the current demand for HPC services far exceeds the current supply at ASC and the other DoD MSRCs. Given the intense demand for these systems, we currently apply the following strategies to meet user expectations and provide the best possible service to the widest possible set of users:

- Limit the number of concurrent batch jobs that run to reduce CPU over-subscription and retain tolerable interactive and multi-processor

job response.

- Relax (just a bit) the desire to have all of central memory in use at any particular instant to be able to schedule and start any batch job that bubbles up to the top of the input queue. Otherwise, some jobs launch regularly and others wait in the input queue for several days.
- Make “time in an input queue” a large factor in selecting jobs for execution. It promotes fairness and encourages users to submit jobs regularly and often to get their fair-share of system resources.
- Encourage small-memory and single processor applications to move to workstations, while retaining the ability for these small applications to execute on our HPC systems.
- Provide special queues and processing times that support special projects such as the DoD Challenge Projects.

## ASC MSRC Batch Job Scheduling Policies

### Interactive Usage Policies

The current interactive usage policy can be found on our website at
http://www.asc.hpc.mil/overall/policy_procedure/policies/use_policy.php

These interactive memory and CPU time limits may be raised on a case-by-case basis by sending a specific request and appropriate justification to the ASC MSRC Service Center.

### Batch Usage Policies

Current queue limits and batch policies can be found on our website at
http://www.asc.hpc.mil/overall/policy_procedure/policies/use_policy.php

### Falsifying Queue Requirements Policy

Users must request ONLY the memory and CPU resources necessary to complete their jobs successfully and efficiently. Users attempting to gain apparent advantage by falsifying memory and/or CPU resource requirements to have their job placed in a queue which apparently has less work waiting to run in it will have thief jobs killed without notice. The net result of such requests is a reduction in overall batch throughput for all users since the queue scheduler honors the excess memory request and does not schedule other work. Jobs found running in the wrong queues because of falsified memory and/or CPU requests are subject to deletion.

## E. Security/Authentication Environment and Policies

### Account Protection

The ASC MSRC uses a combination of Kerberos, SecurID, and Secure Shell (ssh) to prevent unauthorized access to MSRC resources. Kerberos is an authentication system that utilizes a series of encrypted messages sent between two systems to verify the identity of someone attempting access. SecurID is a card based system used to generate a unique passcode each time it is utilized. Secure Shell is a program that allows a user to login over the network, to execute commands on another system, to copy files from one system to another, and to initiate secure X sessions.

### Obtaining an Account at the ASC MSRC

Defense Research, Development, Test and Evaluation users who are interested in obtaining an account on any of the ASC MSRC resources should contact their local S/AAA. In the event you do not know who this person is, please call the ASC MSRC Customer Support at 1-888-677-2272 or e-mail hpc-accounts@asc.hpc.mil. More information is available at http://www.hpcmo.hpc.mil/Htdocs/SAAA.

### Treatment of Passwords

All ASC MSRC systems require a userid and password. Passwords must be protected. You should change the password whenever it is compromised. The system will force a password change every 90 days. ASC MSRC cannot decipher your encrypted password. However, the MSRC can reset your password. **Users are not permitted to share userids on any of the systems**. A user will be permanently deactivated from all the ASC MSRC machines if he has shared his account or password.

### Visitor Information

Detailed visitor information can be found in Section A.

## F. ASC MSRC WWW Information and Site Specific Documentation

### WWW Information

Information about the ASC MSRC can be found on our website:
http://www.asc.hpc.mil

### On-line Documentation

On-line versions of the User's Guides are available on our website:
http://www.asc.hpc.mil/customer/userdocs

## G. ASC MSRC Points of Contact

Wright-Patterson Air Force Base

Public Affairs Office
4375 Chidlaw Road, Suite 6
Wright-Patterson AFB, Ohio 45433-5006
(937) 257-7592
DSN: 787-7592

ASC MSRC Customer Assistance and Technology Center

Toll Free Phone:1 888 MSRC ASC
1-888-677-2272
Commercial Phone:937-255-0194
DSN:785-0194
FAX:937-656-9538
E-mail:msrchelp@asc.hpc.mil